WHAT IS CLAIMED IS:

1. A switching system configured to switch packets of information between ports, each packet comprising a header, each header comprising a destination port address, the switching system comprising:

a plurality of port processors, each port processor being configured to receive and transmit packets of information to and from a line card coupled to each port processor, each port processor being further configured to buffer a plurality of incoming packets from a line card;

a scheduler configured to receive the headers of incoming packets from the port processors, the scheduler being further configured to resolve contentions between headers with destination port addresses that are at least partially similar, the scheduler being further configured to send a request grant to each port processor that sent a header which (a) did not have destination port address that is at least partially similar with destination port addresses of other headers or (b) had a priority level higher than other headers with destination port addresses that are at least partially similar; and

a switch fabric configured to receive incoming packets from the port processors as specified by the request grants from the scheduler and transmit the packets to the port processors as specified by the destination port addresses of the packets, wherein each port processor is configured to transmit packets from the switch fabric to a line card.

2. The switching system of Claim 1, wherein each port processor comprises an input buffer comprising a plurality of sub-queues, the input buffer configured to buffer incoming packets from a line card in the plurality of sub-queues based on the destination port addresses of the packet headers.

3. The switching system of Claim 1, wherein the scheduler comprises a plurality of randomizers, each randomizer configured to receive a plurality of headers from a plurality of port processors, each randomizer comprising:

a memory storing a plurality of randomization permutations;

a pointer coupled to the memory, the pointer being configured to select one randomization permutation at a time; and

a routing crossbar coupled to the memory, the routing crossbar being configured to route sets of headers to a plurality of routing modules coupled to the randomizer according to a randomization permutation selected by the pointer.

4. The switching system of Claim 3, wherein each port processor comprises:

a memory storing a plurality of randomization permutations, the randomization permutations in the port processor being substantially similar to the randomization permutations in each randomizer; and

a pointer coupled to the memory, the pointer being configured to select one randomization permutation at a time, the pointer in the port processor being synchronized with the pointer in the randomizer, wherein each port processor is configured to transfer a packet and a randomization permutation selected by the pointer to a first stage routing module in the switch fabric, the first stage routing module being configured to receive a plurality of packets and one randomization permutation at a time from a plurality of port processors, the first stage routing module being further configured to use one randomization permutation at a time to route the packets to a plurality of second stage routing modules in the switch fabric.

5. The switching system of Claim 4, wherein the first stage routing module in the switch fabric uses the same randomization permutation previously used by the randomizer in the scheduler.

6. The switching system of Claim 1, wherein each port processor comprises an output buffer configured to buffer packets from the switch fabric before transmitting the packets to a line card.

7. The switching system of Claim 1, wherein the scheduler comprises a first stage of routing modules, each routing module of the first stage being coupled to a plurality of port processors, each routing module of the first stage being configured

to (a) receive a plurality of incoming headers from a plurality of port processors; (b) resolve contentions between headers with a similar first set of bits in the destination port addresses of the headers by comparing priority bits in the headers; and (c) output any non-contending headers and one or more contending headers with the highest priority bits on separate output lines.

8. The switching system of Claim 7, wherein the scheduler further comprises a second stage of routing modules, each routing module of the second stage being coupled to a plurality of routing modules of the first stage, each routing module of the second stage being configured to (a) receive a plurality of headers from a plurality of routing modules of the first stage; (b) resolve contentions between headers with a similar second set of bits in the destination port addresses of the headers by comparing priority bits of the headers; and (c) output any non-contending headers and one or more contending headers with the highest priority bits on separate output lines.

9. The switching system of Claim 8, wherein the scheduler further comprises a plurality of scheduler port controllers coupled between the port processors and the first stage, each scheduler port controller comprising a plurality of sub-queues, each scheduler port controller being configured to buffer incoming headers from a port processor in the sub-queues based on the destination port addresses of the headers, each scheduler port controller being further configured to transfer a header from a first sub-queue to a routing module of the first stage at the beginning of a first sub-cycle clock signal, if the scheduler port controller does not receive an acknowledgment signal from the first stage at the end of the first sub-cycle, the scheduler port controller transfers a header from a second sub-queue to the routing module of the first stage at the beginning of a second sub-cycle clock signal.

10. The switching system of Claim 9, wherein the scheduler further comprises acknowledgment crossbars in the first and second stages, each acknowledgment crossbar configured to transmit one or more acknowledgment signals.

11. The switching system of Claim 7, wherein each routing module is further configured to output only headers with asserted valid bits in the headers.

12. The switching system of Claim 7, wherein each routing module comprises:

a routing crossbar coupled to a plurality of port processors, the routing crossbar comprising a plurality of crosspoints configured to connect a plurality of input lines to a plurality of output lines, each input line being coupled to port processor;

a plurality of control units coupled to the crossbar, each control unit being configured to (a) determine whether destination addresses of incoming headers match an address of one of the output lines of the routing crossbar; (b) resolve contentions between headers with destination addresses that match an address of one of the output lines by comparing priority bits of the headers; and (c) activate a crosspoint in the crossbar to output a non-contending header or a contending header with the highest priority bits on one of the output lines of the crossbar; and

an acknowledgment crossbar coupled to the control units, the acknowledgment crossbar comprising a plurality of electronic transistors configured to connect a plurality of input lines to a plurality of output lines, the output lines being coupled to a plurality of port processors, the acknowledgment crossbar being configured to transfer one or more acknowledgment signals to one or more port processors.

13. The switching system of Claim 12, wherein each routing module further comprises an input register coupled between a plurality of port processors and the input lines of the routing crossbar, the input register being configured to buffer a plurality of incoming headers from the plurality of port processors until one or more crosspoints in the crossbar has been activated by one or more control units.

14. The switching system of Claim 12, wherein each control unit comprises:

a plurality of address matching filters, each address matching filter being coupled to one of the input lines of the routing module, each address matching

filter being configured to append a match bit to a header with a destination address that matches a pre-stored value;

a sorter coupled to the plurality of address matching filters, the sorter configured to sort one or more headers output by the address matching filters according to any match bits and priority bits of the headers;

a decoder coupled to the sorter, the decoder being configured to activate a crosspoint in the crossbar to output a header with an asserted match bit and priority bits that are higher than other headers with asserted match bits.

15. The switching system of Claim 14, wherein each control unit further comprises a locking circuit coupled to the decoder, the locking circuit being configured to lock the decoder if a previous header passed through the control unit during a request-grant clock cycle.

16. The switching system of Claim 1, wherein the scheduler is a single integrated chip.

17. The switching system of Claim 1, wherein the switch fabric comprises a first stage of routing modules, each routing module of the first stage being coupled to a plurality of port processors, each routing module of the first stage being configured to (a) receive a plurality of incoming headers from a plurality of port processors; and (b) route headers according to the destination addresses of the headers.

18. The switching system of Claim 1, wherein the switch fabric comprises a second stage of routing modules, each routing module of the second stage being coupled to a plurality of routing modules of the first stage, each routing module of the second stage being configured to (a) receive a plurality of incoming headers from a plurality of the routing modules of the first stage; and (b) route headers to the port processors according to the destination addresses of the headers.

19. The switching system of Claim 1, wherein the scheduler is configured to receive 64 headers from the port processors, resolve one or more contentions

between the headers in 20 nanoseconds or less, and output about 16 request grants to the port processors.

20. The switching system of Claim 1, wherein the scheduler is configured to receive 128 headers from the port processors, resolve one or more contentions between the headers in 20 nanoseconds or less, and output about 16 request grants to the port processors.

21. The switching system of Claim 1, wherein the scheduler is configured to receive headers from 32 port processors, resolve one or more contentions between the headers, and output request grants two or more of the port processors.

22. The switching system of Claim 1, wherein the scheduler is configured to receive headers from 64 or more port processors, resolve one or more contentions between the headers, and output request grants two or more of the port processors.

23. A scheduler in a switching system configured to switch packets of information between ports, each packet comprising a header, each header comprising a destination port address and a priority level, the scheduler comprising:

a plurality of randomizers, each randomizer being coupled to a set of port processors, each randomizer being configured to receive headers from the set of port processors and randomly output the headers;

a first stage of routing modules, each routing module in the first stage being coupled to at least two randomizers, each routing module in the first stage being configured to receive headers from at least two randomizers, each routing module in the first stage being configured to output one or more headers with a destination address that matches a pre-stored value and have a higher priority level than other headers with destination addresses that match the pre-stored value; and

a second stage of routing modules, each routing module in the second stage being coupled to at least two routing modules in the first stage, each routing module in the second stage being configured to receive headers from at least two routing modules in the first stage, each routing module in the second stage being configured to output an acknowledgment signal to each port processors that sent a

header with a destination address that matches a pre-stored value and has a higher priority level than other headers with destination addresses that match the pre-stored value.

24. The scheduler of Claim 23, further comprising a plurality of scheduler port controllers, each scheduler port controller being coupled between a port processor and a randomizer, each scheduler port controller being configured to receive a plurality of headers from a port processor.

25. A method of switching packets of information between ports, each packet comprising a header, each header comprising a destination port address and a priority level, the method comprising:

buffering a plurality of incoming packets;

resolving contentions between headers with destination addresses that are at least partially similar by selecting headers with a higher priority level than other headers with at least partially similar destination addresses;

sending one or more request grants to one or more port processors, each request grant representing a header that (a) did not contend with other headers or (b) had a higher priority level than other headers with at least partially similar destination addresses; and

switching one or more packets between the port processors based on the one or more request grants.

26. The method of Claim 25, wherein buffering a plurality of incoming packets comprises storing incoming packets in a plurality of sub-queues in a plurality of scheduler port controllers, each sub-queue storing packets with a particular destination address, each scheduler port controller sending a first packet from a first sub-queue to a scheduler switch fabric to resolve contentions, if the first packet is unsuccessful, each scheduler port controller sends a second packet from a second sub-queue to the scheduler switch fabric to resolve contentions.

27. The method of Claim 25, wherein resolving contentions comprises:

randomly routing headers of packets from a plurality of port processors to a routing stage;

in the routing stage, determining whether two or more headers have destination addresses that are at least partially similar;

from the two or more headers with destination addresses that are at least partially similar, selecting a header with a higher priority level; and

sending an acknowledgment signal to each port processor that sent a header that (a) did not contend with other headers or (b) had a higher priority level than other headers with at least partially similar destination addresses.

28. The method of Claim 27, further comprising locking each control unit in the routing stage that output a header that did not contend with other headers or had a higher priority level than other headers with at least partially similar destination addresses.